arXiv:1607.07967v2 [cs.DB] 7 Oct 2016

PIWD: A Plugin-based Framework for Well-Designed SPARQL

Xiaowang Zhang[1,3,4], Zhenyu Song[1,3], Zhiyong Feng[2,3], and Xin Wang[1,3]

[1] School of Computer Science and Technology, Tianjin University, Tianjin, China
[2] School of Computer Software, Tianjin University, Tianjin 300350, China
[3] Tianjin Key Laboratory of Cognitive Computing and Application, Tianjin, China
[4] Key Laboratory of Computer Network and Information Integration (Southeast University), Ministry of Education, Nanjing 211189, China

Abstract. In the real world datasets (e.g.,DBpedia query log), queries built on well-designed patterns containing only AND and OPT operators (for short, WDAO-patterns) account for a large proportion among all SPARQL queries. In this paper, we present a plugin-based framework for all SELECT queries built on WDAO-patterns, named PIWD. The framework is based on a parse tree called *well-designed AND-OPT tree* (for short, WDAO-tree) whose leaves are basic graph patterns (BGP) and inner nodes are the OPT operators. We prove that for any WDAO-pattern, its parse tree can be equivalently transformed into a WDAO-tree. Based on the proposed framework, we can employ any query engine to evaluate BGP for evaluating queries built on WDAO-patterns in a convenient way. Theoretically, we can reduce the query evaluation of WDAO-patterns to subgraph homomorphism as well as BGP since the query evaluation of BGP is equivalent to subgraph homomorphism. Finally, our preliminary experiments on gStore and RDF-3X show that PIWD can answer all queries built on WDAO-patterns effectively and efficiently.

Keywords: SPARQL, BGP, well-designed patterns, subgraph homomorphism

1 Introduction

Resource Description Framework (RDF) [30] is the standard data model in the semantic web. RDF describes the relationship of entities or resources using directed labelling graph. RDF has a broad range of applications in the semantic web, social network, bio-informatics, geographical data, etc. [3]. The standard query language for RDF graphs is SPARQL [24]. Though SPARQL is powerful to express queries over RDF graphs [2], generally, the query evaluation of the full SPARQL is PSPACE-complete [22].

Currently, there are some popular query engines for supporting the full SPARQL such as Jena [8] and Sesame [7]. However, they become not highly efficient when they handle some large RDF datasets[40,41]. Currently, gStore[40,41] and RDF-3X[20] can highly efficiently query large datasets. But gStore and

RDF-3X merely provide querying services of BGP. Therefore, it is very necessary to develop a query engine with supporting more expressive queries for large datasets.

Since the OPT operator is the least conventional operator among SPARQL operators [37], it is interesting to investigate those patterns extending BGP with the OPT operator. Let us take a look at the following example.

An RDF example in Table 1 describes the entities of bloggers and blogs. The relationship between a blogger and a blog is revealed in the property of *foaf:maker*. Both blogger and blog have some properties to describe themselves. Triples can be modeled as a directed graph substantially.

Table 1. bloggers.rdf

Subject	Predict	Object
id1	foaf:name	Jon Foobar
id1	rdf:type	foaf:Agent
id1	foaf:weblog	foobar.xx/blog
foobar.xx/blog	dc:title	title
foobar.xx/blog	rdfs:seeAlso	foobar.xx/blog.rdf
foobar.xx/blog.rdf	foaf:maker	id1
foobar.xx/blog.rdf	rdf:type	rss:channel

Example 1. Consider the RDF dataset G storing information in Table 1. Given a BGP $Q = ((?x, \textit{foaf:maker}, ?y)$ AND $(?z, \textit{foaf:name}, ?u))$, its evaluation over G is as follows:

$[\![Q]\!]_G =$

?x	?y	?z	?u
foobar.xx/blog.rdf	id1	id1	Jon Foobar

Consider a new pattern Q_1 obtained from Q by adding the OPT operator in the following way:
$Q_1 = (((?x, \textit{foaf:maker}, ?y)$ OPT $(?y, \textit{rdf:type}, ?v))$ AND $(?z, \textit{foaf:name}, ?u))$, the evaluation of Q_1 over G is as follows:

$[\![Q_1]\!]_G =$

?x	?y	?v	?z	?u
foobar.xx/blog.rdf	id1	foaf:Agent	id1	Jon Foobar

Consider another pattern $Q_2 = (((?x, \textit{foaf:maker}, ?y)$ OPT $(?y, \textit{rdf:type}, ?z))$ AND $(?z, \textit{foaf:name}, ?u))$, the evaluation of Q_2 over G is the empty set, i.e., $[\![Q_2]\!]_G = \emptyset$.

In the above example, Q_1 is a well-designed pattern while Q_2 is not a well-designed pattern [22].

In fact, we investigate that queries built on well-designed patterns are very popular in a real world. For example, in LSQ[26], a Linked Dataset describing SPARQL queries extracted from the logs of four prominent public SPARQL endpoints containing more than one million available queries shown in Table 2, queries built on well-designed patterns are over 70% [12,29].

Table 2. SPARQL logs source in LSQ

Dataset	Date	Triple Number
DBpedia	30/04/2010 to 20/07/2010	232,000,000
Linked Geo Data (LGD)	24/11/2010 to 06/07/2011	1,000,000,000
Semantic Web Dog Food (SWDF)	16/05/2014 to 12/11/2014	300,000
British Museum (BM)	08/11/2014 to 01/12/2014	1,400,000

Furthermore, queries with well-designed AND-OPT patterns (for short, WDAO-patterns) are over 99% among all queries with well-designed patterns in LSQ [12,29]. In short, the fragment of WDAO-patterns is a natural extension of BGP in our real world. Therefore, we mainly discuss WDAO-patterns in this paper.

In this paper, we present a plugin-based framework for all SELECT queries built on WDAO-patterns, named PIWD. Within this framework, we can employ any query engine evaluating BGP for evaluating queries built on WDAO-patterns in a convenient way. The main contributions of this paper can be summarized as follows:

- We present a parse tree named *well-designed AND-OPT tree* (for short, WDAO-tree), whose leaves are BGP and all inner nodes are the OPT operator and then prove that for any WDAO-pattern, it can be translated into a WDAO-tree.
- We propose a plugin-based framework named *PIWD* for query evaluation of queries built on WDAO-patterns based on WDAO-tree. Within this framework, a query could be evaluated in the following three steps: (1) translating that query into a WDAO tree T; (2) evaluating all leaves of T via query engines of BGP; and (3) joining all solutions of children to obtain solutions of their parent up to the root.
- We implement the proposed framework PIWD by employing gStore and RDF-3X and evaluate the experiments on LUBM.

The rest of this paper is organized as follows: Section 2 briefly introduces the SPARQL, conception of well-designed patterns and OPT normal form. Section 3 defines the well-designed and-opt tree to capture WDAO-patterns. Section 4 presents PIWD and Section 5 evaluates experimental results. Section 6 summarizes our related works. Finally, Section 7 summarizes this paper.

2 Preliminaries

In this section, we introduce RDF and SPARQL patterns, well-designed patterns, and OPT normal form [22].

2.1 RDF

Let I, B and L be infinite sets of *IRIs*, *blank nodes* and *literals*, respectively. These three sets are pairwise disjoint. We denote the union $I \cup B \cup L$ by U, and elements of $I \cup L$ will be referred to as *constants*.

A triple $(s, p, o) \in (I \cup B) \times I \times (I \cup B \cup L)$ is called an *RDF triple*. A *basic graph pattern* (BGP) is a set of triple patterns.

2.2 Semantics of SPARQL patterns

The semantics of patterns is defined in terms of sets of so-called *mappings*, which are simply total functions $\mu\colon S \to U$ on some finite set S of variables. We denote the domain S of μ by $\mathrm{dom}(\mu)$.

Now given a graph G and a pattern P, we define the semantics of P on G, denoted by $[\![P]\!]_G$, as a set of mappings, in the following manner.

- If P is a triple pattern (u, v, w), then
$$[\![P]\!]_G := \{\mu\colon \{u, v, w\} \cap V \to U \mid (\mu(u), \mu(v), \mu(w)) \in G\}.$$
Here, for any mapping μ and any constant $c \in I \cup L$, we agree that $\mu(c)$ equals c itself. In other words, mappings are extended to constants according to the identity mapping.
- If P is of the form P_1 UNION P_2, then $[\![P]\!]_G := [\![P_1]\!]_G \cup [\![P_2]\!]_G$.
- If P is of the form P_1 AND P_2, then $[\![P]\!]_G := [\![P_1]\!]_G \bowtie [\![P_2]\!]_G$, where, for any two sets of mappings Ω_1 and Ω_2, we define
$$\Omega_1 \bowtie \Omega = \{\mu_1 \cup \mu_2 \mid \mu_1 \in \Omega_1 \text{ and } \mu_2 \in \Omega_2 \text{ and } \mu_1 \sim \mu_2\}.$$
Here, two mappings μ_1 and μ_2 are called *compatible*, denoted by $\mu_1 \sim \mu_2$, if they agree on the intersection of their domains, i.e., if for every variable $?x \in \mathrm{dom}(\mu_1) \cap \mathrm{dom}(\mu_2)$, we have $\mu_1(?x) = \mu_2(?x)$. Note that when μ_1 and μ_2 are compatible, their union $\mu_1 \cup \mu_2$ is a well-defined mapping; this property is used in the formal definition above.
- If P is of the form P_1 OPT P_2, then
$$[\![P]\!]_G := ([\![P_1]\!]_G \bowtie [\![P_2]\!]_G) \cup ([\![P_1]\!]_G \setminus [\![P_2]\!]_G),$$
where, for any two sets of mappings Ω_1 and Ω_2, we define
$$\Omega_1 \setminus \Omega_2 = \{\mu_1 \in \Omega_1 \mid \neg\exists \mu_2 \in \Omega_2 : \mu_1 \sim \mu_2\}.$$
- If P is of the form $\mathrm{SELECT}_S(P_1)$, then $[\![P]\!]_G = \{\mu|_{S \cap \mathrm{dom}(\mu)} \mid \mu \in [\![P_1]\!]_G\}$, where $f|_X$ denotes the standard mathematical notion of restriction of a function f to a subset X of its domain.
- Finally, if P is of the form P_1 FILTER C, then $[\![P]\!]_G := \{\mu \in [\![P_1]\!]_G \mid \mu(C) = \mathit{true}\}$.
Here, for any mapping μ and constraint C, the evaluation of C on μ, denoted by $\mu(C)$, is defined in terms of a three-valued logic with truth values *true*, *false*, and *error*. Recall that C is a boolean combination of atomic constraints.
For a bound constraint $\mathrm{bound}(?x)$, we define:
$$\mu(\mathrm{bound}(?x)) = \begin{cases} \mathit{true} & \text{if } ?x \in \mathrm{dom}(\mu); \\ \mathit{false} & \text{otherwise.} \end{cases}$$
For an equality constraint $?x = ?y$, we define:
$$\mu(?x = ?y) = \begin{cases} \mathit{true} & \text{if } ?x, ?y \in \mathrm{dom}(\mu) \text{ and } \mu(?x) = \mu(?y); \\ \mathit{false} & \text{if } ?x, ?y \in \mathrm{dom}(\mu) \text{ and } \mu(?x) \neq \mu(?y); \\ \mathit{error} & \text{otherwise.} \end{cases}$$

Thus, when $?x$ and $?y$ do not both belong to $\mathrm{dom}(\mu)$, the equality constraint evaluates to *error*. Similarly, for a constant-equality constraint $?x = c$, we define:

$$\mu(?x = c) = \begin{cases} \mathit{true} & \text{if } ?x \in \mathrm{dom}(\mu) \text{ and } \mu(?x) = c; \\ \mathit{false} & \text{if } ?x \in \mathrm{dom}(\mu) \text{ and } \mu(?x) \neq c; \\ \mathit{error} & \text{otherwise.} \end{cases}$$

A boolean combination is then evaluated using the truth tables given in Table 3.

Table 3. Truth tables for the three-valued semantics.

p	q	$p \wedge q$	$p \vee q$
true	*true*	*true*	*true*
true	*false*	*false*	*true*
true	*error*	*error*	*true*
false	*true*	*false*	*true*
false	*false*	*false*	*false*
false	*error*	*false*	*error*
error	*true*	*error*	*true*
error	*false*	*false*	*error*
error	*error*	*error*	*error*

p	$\neg p$
true	*false*
false	*true*
error	*error*

2.3 Well-Designed Pattern

A UNION-*free* pattern P is *well-designed* if the followings hold:

- P is safe;
- for every subpattern Q of form $(Q_1 \text{ OPT } Q_2)$ of P and for every variable $?x$ occurring in P, the following condition holds:

  If $?x$ occurs both inside Q_2 and outside Q, then it also occurs in Q_1.

Consider the definition of well-designed patterns, some conceptions can be explained as follows:

Remark 1. In the fragment of and-opt patterns, we exclude FILTER and UNION operators and it contains only AND and OPT operators at most. It is obvious that *and-opt* pattern must be UNION-*free* and safe.

We can conclude that WDAO-patterns are decided by variables in subpattern.

- **UNION-*free* Pattern**: P is UNION-free if P is constructed by using only operators AND, OPT, and FILTER. Every graph pattern P is equivalent to a pattern of the form denoted by (P_1 UNION P_2 UNION $\cdots$ UNION P_n). Each P_i ($1 \leq i \leq n$) is UNION-free.

– **Safe** : If the form of (P FILTER R) holds the condition of $var(R) \subseteq var(P)$, then it is safe.

Note that the OPT operator provides really optional left-outer join due to the weak monotonicity [22]. A SPARQL pattern P is said to be weakly monotone if for every pair of RDF graphs G_1, G_2 such that $G_1 \subseteq G_2$, it holds that $[\![P]\!]_{G_1} \sqsubseteq [\![P]\!]_{G_2}$. In other words, we assume μ_1 represents $[\![P]\!]_{G_1}$, and μ_2 represents $[\![P]\!]_{G_2}$. Then there exists μ' such that $\mu_2 = \mu_1 \cup \mu'$. Weakly monotone is an important property to characterize the satisfiability of SPARQL [38]. For instance, consider the pattern Q_1 in Section 1, $(?y, \textit{rdf:type}, ?v)$ are really optional.

2.4 OPT Normal Form

A UNION-free pattern P is in *OPT normal form* [22] if P meets one of the following two conditions:

– P is constructed by using only the AND and FILTER operators;
– $P = (P_1 \text{ OPT } P_2)$ where P_1 and P_2 patterns are in OPT normal form.

For instance, the pattern Q aforementioned in Section 1 is in OPT normal form. However, consider the pattern $(((?x, p, ?y) \text{ OPT } (?x, q, ?z))\text{AND}(?x, r, ?z))$ is not in OPT normal form.

3 Well-Designed And-Opt Tree

In this section, we propose the conception of the well-designed and-opt tree (WDAO-tree), any WDAO-pattern can be seen as an WDAO-tree.

3.1 WDAO-tree Structure

Definition 1 (WDAO-tree). *Let P be a well-designed pattern in OPT normal form. A well-designed tree T based on P is a redesigned parse tree, which can be defined as follows:*

– *All inner nodes in T are labeled by the* OPT *operator and leaves are labeled by BGP.*
– *For each subpattern $(P_1 \text{ OPT } P_2)$ of P, the well-designed tree T_1 of P_1 and the well-designed tree T_2 of P_2 have the same parent node.*

For instance, consider a WDAO-pattern P[5]

$$P = (((p_1 \text{ AND } p_3) \text{ OPT}_2 p_2) \text{ OPT}_1 ((p_4 \text{ OPT}_4 p_5) \text{ OPT}_5 (p_6 \text{ OPT}_6 p_7))).$$

The WDAO-tree T is shown in Figure 1. As shown in this example, BGP - $(p_1 \text{ AND } p_3)$ is the exact matching in P, which corresponds to the non-optional pattern. Besides, in WDAO-tree, it is the leftmost leaf in T. We can conclude that the leftmost node in WDAO-tree means the exact matching in well-designed SPARQL query pattern.

[5] We give each OPT operator a subscript to differentiate them so that readers understand clearly.

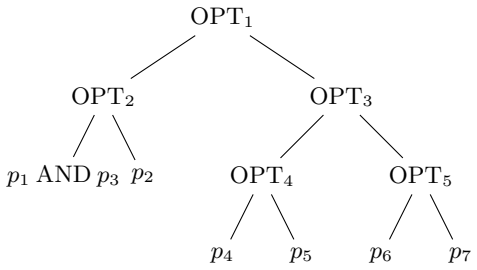

Fig. 1. WDAO-tree

3.2 Rewritting rules over WDAO-tree

As described in Section 1, WDAO-tree does not contain any OPT operator in its leaves. In this sense, patterns as the form of Q_1 in Section 1 cannot be transformed into WDAO-tree since it is not OPT normal form.

Proposition 1. *[22, Theorem 4.11] For every UNION-free well-designed pattern P, there exists a pattern Q in OPT normal form such that P and Q are equivalent.*

In the proof of Proposition 1, we apply three rewriting rules based on the following equations: let P, Q, R be patterns and C a constraint,

- $(P \text{ OPT } R) \text{ AND } Q \equiv (P \text{ AND } Q) \text{ OPT } R$;
- $P \text{ AND } (Q \text{ OPT } R) \equiv (P \text{ AND } Q) \text{ OPT } R$;
- $(P \text{ OPT } R) \text{ FILTER } C \equiv (P \text{ FILTER } C) \text{ OPT } R$.

Intuitively, this lemma states that AND operator can forward and OPT operator can backward in a well-designed pattern with preserving the semantics. The above three rules can be deployed on a WDAO-tree. For each WDAO-tree T, there exists T' corresponding to T after applying rewriting rules.

Figure 2 and Figure 3 have shown that the process of rewriting rules after generating grammar tree and finally WDAO-tree can be obtained. Clearly, WDAO-tree has less height than the grammar tree.

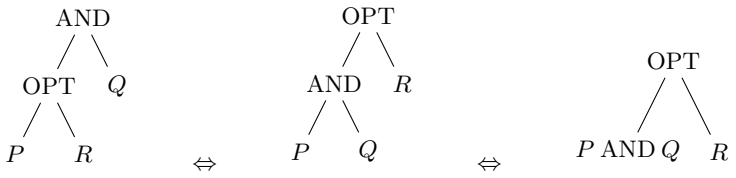

Fig. 2. rewritting rule-1

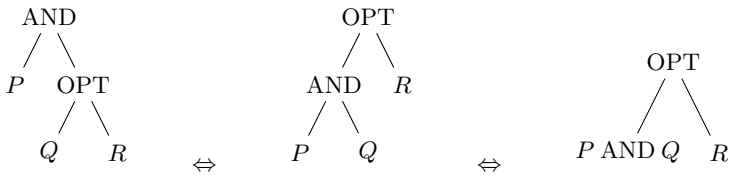

Fig. 3. rewritting rule-2

3.3 WDAO-tree Construction

Before constructing WDAO-tree, we recognize query patterns and attachments at first. Then we rewrite query patterns by rewritting rules, which leads to a new pattern. Based on this new pattern, we construct WDAO-tree in the principle of Definition 1.

In the process of the WDAO-tree construction, we firstly build the grammar tree of SPARQL patterns, whose inner node is either AND operator or OPT operator. This process is based on recursively putting the left pattern and right pattern of operator in the left node and right node respectively until the pattern does not contain any operator. Then we apply the rewritting rules to the grammar tree in Algorithm 1 to build rewriting-tree whose only leaf node is single triple pattern. Different rewritting rules are adopted depending on OPT operator are AND operator's left child or right child. Since WDAO-tree's inner nodes only contain AND operators, After getting rewriting-tree, we merge the AND operators only containing leaf child nodes with its child nodes into new nodes in order to get a WDAO-tree.

The WDAO-tree construction can be executed in PTIME. Given a pattern containing n ANDs and m OPTs, the construction of the grammar tree and rewriting tree have $O(n + m)$ time complexity and $O(nm)$ time complexity, respectively. Furthermore, the merge of nodes whose parent is AND has $O(n)$ time complexity.

4 PIWD Demonstration

In this section, we introduce PIWD, which is a plugin-based framework for well-designed SPARQL.

4.1 PIWD Overview

PIWD is written in Java in a 2-tier design shown in Figure 4. The bottom layer consists of any BGP query framework which is used as a black box for evaluating BGPs. Before answering SPARQL queries, the second layer provides the rewriting process and left-outer join evaluation, which lead to the solutions.

BGP query framework supports both query and RDF data management, such as gStore, RDF-3X and so on, which solve the problem of subgraph isomorphism.

Algorithm 1 rewritting rules

Input: GrammarTree with *Root*;
Output: RewriteTree with *Root*;
while not all AND.child IS OPT **do**
Procedure ReWriteRules(*Root*)
if *Root* IS AND **then**
    if *Root* IS OPT **then**
        swap($Root.left$,$Root.right.left$);
        swap($Root.right$,$Root.right$);
        swap($Root.left.left$,$Root$);
        swap($Root.left.right$,$Root.left.right$);
    end if
    if *Root.right* IS OPT **then**
        swap($Root.left$,$Root.left$);
        swap($Root.right$,$Root.left.left$);
        swap($Root.left.left$,$Root$);
        swap($Root.left.right$,$Root.left.right$);
    end if
end if
Procedure ReWriteRules($Root.left$)
End Procedure
Procedure ReWriteRules($Root.right$)
End Procedure
End Procedure
end while
return *Root*;

PIWD provides the left-outer join between the BGPs. That is, the problem of answering well-designed SPARQL has been transformed into the problem of subgraph isomorphism and left-outer join between triple patterns.

4.2 Answering Queries over PIWD

The query process over PIWD can be described as follows:

Firstly, WDAO-tree is built after rewriting rules on the grammar tree. Secondly, post-order traversal is applied on WDAO-trees. The traversal rule is: If the node is a leaf node without the OPT operator, BGP query framework is deployed on it to answer this query and return solutions which is stored in a stack. If the node is an inner node labeled by the OPT operator, we get the top two elements in the stack and left-outer join them. We repeat this process until all of WDAO-tree's nodes are visited. Finally, only one element in the stack is the final solutions.

In the querying processing, BGP query framework serves as a query engine to support queries from leaves in WDAO-trees. OPT operators take an essential position in the query processing. Users receive optional solutions based on OPT operators which contribute to the semantic abundance degree since optional

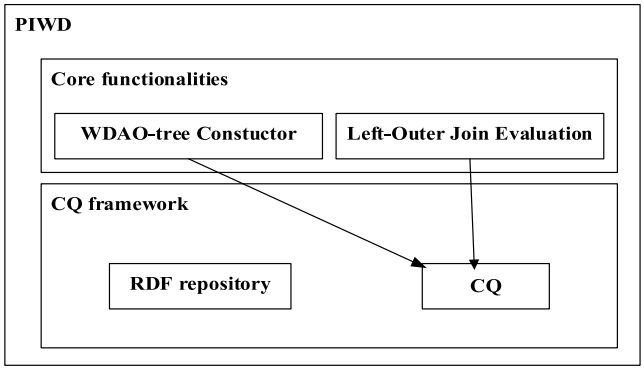

Fig. 4. PIWD architecture

solutions are considered in this sense. In other words, OPT operators lead to the explosive growth of the solution scale.

The query process is described in Algorithm 2.

5 Experiments and Evaluations

This section presents our experiments. The purpose of the experiments is to evaluate the performance of different WDAO-patterns.

5.1 Experiments

Implementations and running environment All experiments were carried out on a machine running Linux, which has one CPU with four cores of 2.40GHz, 32GB memory and 500GB disk storage. All of the algorithms were implemented in Java. gStore[40,41] and RDF-3X[20] are used as the underlying query engines to handle BGPs. In our experiments, there is no optimization in our OPT operation.

gStore and RDF-3X Both gStore and RDF-3X are SPARQL query engines for subgraph matching. gStore stores RDF data in disk-based adjacency lists, whose format is *[vID,vLabel,adjList]*, where *vID* is the vertex ID, *vLabel* is the corresponding URI, and *adjList* is the list of its outgoing edges and the corresponding neighbor vertices. gStore converts an RDF graph into a data signature graph by encoding each entity and class vertex. Some different hash functions such as BKDR and AP hash functions are employed to generate signatures, which compose a novel index (called VS*-tree). A filtering rule and efficient search algorithms are developed for subgraph queries over the data signature graph in order to speed up query processing. gStore can answer exact SPARQL queries and queries with wildcards in a uniform manner. RDF-3X engine is a RISC-style architecture for executing SPARQL queries over large repositories of RDF triple.

Algorithm 2 Query Processing over PIWD

Input: WDAO-tree with *Root*; Prefix $prefix$; *Stack* to store subresults;
Output: Query result *result*;

```
1: Procedure TraverseTree(Root)
2: if root is not null then
3:     Procedure TraverseTree(Root → Lnode)
4:     End Procedure
5:     Procedure TraverseTree(Root → Rnode)
6:     End Procedure
7:     if node is not OPTIONAL then
8:         subquery=AssembleQuery(prefix,node);
9:         subresult=QueryIngStore(subquery);
10:        Push(Stack , subresult);
11:    else
12:        r=Pop(Stack);
13:        l=Pop(Stack);
14:        result=l ⟗ r;
15:        Push(Stack , result);
16:    end if
17: end if
18: End Procedure
19: list=ConvertToList(Stack);
20: return list;
```

Physical design is workload-independent by creating appropriate indexes over a single giant triples table in RDF-3X. And the query processor is RISC-style by relying mostly on merge joins over sorted index lists. gStore and RDF-3X have good performances in BGPs since their query methods are based on subgraph matching.

Dataset We used LUBM[6] as the dataset in our experiments to investigate the relationship between query response time and dataset scale. LUBM, which features an ontology for the university domain, is a standard benchmark to evaluate the performance of semantic Web repositories, In our experiments, we used LUBM1, LUBM50, LUBM100, LUBM150 and LUBM200 as our query datasets. The LUBM dataset details in our experiments are shown in Table 4.

SPARQL queries The queries over LUBM were designed as four different forms, which corresponds to different WDAO-trees. The details of queries are described in Table 5. Clearly, OPT nesting in Q_2 is the most complex among four forms. Furthermore, we build the AND operator in each query.

5.2 Evaluation on PIWD

The variation tendencies of query response time are shown in Table 6, Table 7 and Figure 5. Query efficiency is decreased with higher response time when

[6] http://swat.cse.lehigh.edu/projects/lubm/

Table 4. LUBM Dataset Details

Dataset	Number of triples	RDF NT File Size(bytes)
LUBM1	103,104	14,497,954
LUBM50	6,890,640	979,093,554
LUBM100	13,879,971	1,974,277,612
LUBM150	20,659,276	2,949,441,119
LUBM200	27,643,644	3,954,351,227

Table 5. SPARQL queries Details

QueryID	Pattern	OPT amount
Q_1	$(P_1$ AND P_2 AND $P_3)$ OPT P_4	1
Q_2	$((P_1$ AND P_2 AND $P_3)$ OPT $P_4)$ OPT $(P_5$ OPT $P_6)$	3
Q_3	$((P_1$ AND P_2 AND $P_3)$ OPT $P_4)$ OPT P_5	2
Q_4	P_1 OPT $((P_2$ AND P_3 AND $P_4)$ OPT $P_5)$	2

OPT nesting becomes more complex. Furthermore, there has been a significant increase in query response time when the dataset scale grows up. For instance, we observe Q_2, which corresponds to the most complex pattern in our four experimental SPARQL patterns. When the dataset is ranging from LUBM100 to LUBM200, its query response time extends more than five times even though the dataset scale extends two times. In this sense, OPT nesting complexity in WDAO-patterns influences query response time especially for large dataset scale.

Table 6. Query Response Time[ms] on gStore

	LUBM1	LUBM50	LUBM100	LUBM150	LUBM200
Q_1	1,101	617,642	1,329,365	2,126,383	2,978,237
Q_2	1,870	1,010,965	2,901,295	6,623,806	10,041,836
Q_3	1,478	637,128	1,359,315	2,191,356	3,068,692
Q_4	1,242	644,155	1,456,232	2,151,811	3,129,246

6 Related works

In this section, we survey related works in the following three areas: BGP query evaluation algorithms, well-designed SPARQL and BGP query evaluation frameworks.

BGP query algorithms have been developed for many years. Existing algorithms mainly focus on finding all embedding in a single large graph, such as ULLmann[31], VF2[18], QUICKSI[28], GraphQL[14], SPath[39], STW[32] and TurboIso[13]. Some optimization method has been adapted in these techniques, such as adjusting matching order, pruning out the candidate vertices. However,

Table 7. Query Response Time[ms] on RDF-3X

	LUBM1	LUBM50	LUBM100	LUBM150	LUBM200
Q_1	1,231	625,703	1,401,782	2,683,461	3,496,156
Q_2	1,900	1,245,241	2,983,394	7,286,812	10,852,761
Q_3	1,499	640,392	1,427,392	2,703,981	3,672,970
Q_4	1,316	648,825	1,531,547	2,791,152	3,714,042

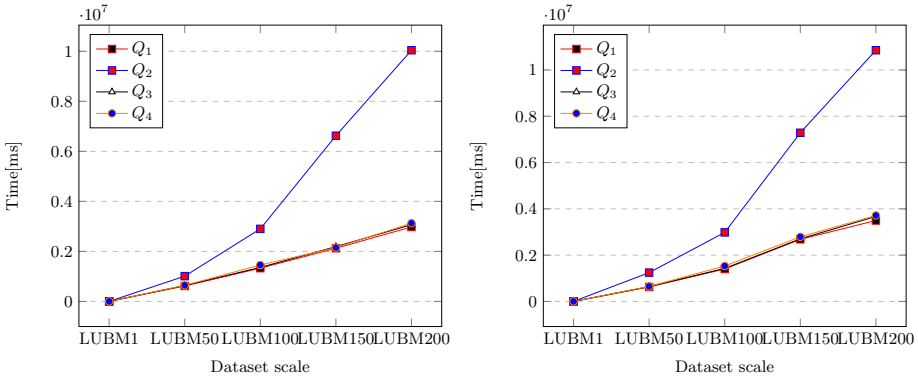

(a) Performance on gStore (b) Performance on RDF-3X

Fig. 5. Query Response Time over LUBM

the evaluation of well-designed SPARQL is not equivalent to the BGP query evaluation problem since there exists inexact matching.

It has been shown that the complexity of the evaluation problem for the well-designed fragment is coNP-complete[22]. The quasi well-designed pattern trees (QWDPTs), which are undirected and ordered, has been proposed [16]. This work aims at the analysis of containment and equivalence of well-designed pattern. Efficient evaluation and semantic optimization of WDPT have been proposed in [5]. Sparm is a tool for SPARQL analysis and manipulation in [17]. Above-mentioned all aim at checking well-designed patterns or complexity analysis without evaluation on well-designed patterns. Our WDAO-tree is different from QWDPTs in structure and it emphasizes reconstructing query plans. The OPT operation optimization has been proposed in [19], which is different from our work since our work aims to handle a plugin in any BGP query engine in order to deal with WDAO-patterns in SPARQL queries.

RDF-3X[20], TripleBit[34], SW-Store[1], Hexastore[33] and gStore[40,41] have high performance in BGPs. RDF-3X create indexes in the form of B+ tree, as well as TripleBit in the form of ID-Chunk. All of them have efficient performance since they concentrate on the design of indexing or storage. However, they can only support exact SPARQL queries, since they replace all literals (in RDF triples) by ids using a mapping dictionary. In other words, they cannot support WDAO-patterns well. Virtuoso[9] and MonetDB[6] support open-source

and commercial services. Jena[8] and Sesame[7] are free open source Java frameworks for building semantic web and Linked Data applications, which focus on SPARQL parse without supporting large-scale date. Our work is independent on these BGP query frameworks, and any BGP query engine is adaptable for our plugin.

7 Conclusion

In this paper, we have presented PIWD, which is a plugin adaptable for any BGP query framework to handle WDAO-patterns. Theoretically, PIWD rebuilds the query evaluation plan based on WDAO-trees. After employing BGP query framework on WDAO-trees, PIWD supports the left-outer join operation between triple patterns. Our experiments show that PIWD can deal with complex and multi-level nested WDAO-patterns. In the future, we will further handle other non-well-designed patterns and deal with more operations such as UNION. Besides, we will consider OPT operation optimization to improve efficiency of PIWD and implement our framework on distributed RDF graphs by applying the distributed gStore [21].

Acknowledgments

This work is supported by the programs of the National Key Research and Development Program of China (2016YFB1000603), the National Natural Science Foundation of China (NSFC) (61502336), and the open funding project of Key Laboratory of Computer Network and Information Integration (Southeast University), Ministry of Education (K93-9-2016-05). Xiaowang Zhang is supported by Tianjin Thousand Young Talents Program and the project-sponsored by School of Computer Science and Technology in Tianjin University.

References

1. Abadi D.J., Marcus A., Madden S.R., and Hollenbach K. (2009). SW-Store: A vertically partitioned dbms for semantic web data management. *VLDB J.*, 18(2):385–406.
2. Angles R. and Gutierrez C.(2008). The expressive power of SPARQL. In: *Proc. of ISWC 2008*, pp. 114–129.
3. Arenas M. and Pérez J. (2011). Querying semantic web data with SPARQL. In: *Proc. of SIGMOD 2011*, pp.305–316.
4. Auer S., Bizer C., Kobilarov G., Lehmann J., Cyganiak R., and Ives Z.(2007). DBpedia: A nucleus for a web of open data. In: *Proc. of ISWC 2007*, pp. 11–15.
5. Barcelo P., Pichler R., and Skritek S. (2015). Efficient evaluation and approximation of well-designed pattern trees. In: *Proc. of PODS 2015*, pp.131–144.
6. Boncz P.A., Zukowski M., and Nes N.J. (2005). MonetDB/x100: Hyper-pipelining query execution. In: *Proc. of CIDR 2005*.

7. Broekstra J., Kampman A., and Van Harmelen F. (2002). *Sesame: A generic architecture for storing and querying RDF and RDF schema.* In: *Proc. of ISWC 2002*, pp.54–68.
8. Carroll J.J., Dickinson I., Dollin C., Reynolds D., Seaborne A., and Wilkinson K. (2004). Jena: Implementing the semantic web recommendations. In: *Proc. of WWW 2004*, pages 74–83.
9. Erling O. and Mikhailov I. (2009). RDF support in the virtuoso DBMS. *Studies in Computational Intelligence*, pp. 7–24.
10. Green T.J. (2011). Containment of conjunctive queries on annotated relations. *Theory Comput. Syst.*, 49(2): 429–459.
11. Garey M.R. and Johnson D.S. (1979). Computers and intractability: A guide to the theory of NP-completeness. *W.H. Freeman and Company.*
12. Han X., Feng Z., Zhang X., Wang X., and Rao G. (2016). On the statistical analysis of practical SPARQL queries. In: *Proc. of WebDB 2016*, article 2.
13. Han W.S., Lee J., and Lee J. H.(2013). Turbo iso: Towards ultrafast and robust subgraph isomorphism search in large graph databases. In: *Proc. of SIGMOD 2013*, pp. 337–348.
14. He H. and Singh A.K. (2010). *Query language and access methods for graph databases.* Springer.
15. Kostylev E.V. and Grau B.C. (2014). *On the semantics of SPARQL queries with optional matching under entailment regimes.* In: *Proc. of ISWC 2014*, pp. 374–389.
16. Letelier A., Pérez J., Pichler R., and Skritek S.(2012). Static analysis and optimization of semantic web queries. In: *Proc. of PODS 2012*, 38(4):84–87.
17. Letelier A., Pérez J., Pichler R., and Skritek S. (2012). SPAM: A SPARQL analysis and manipulation tool. In: *Proc. of VLDB 2012*, 5(12):1958–1961.
18. Luigi P C., Pasquale F., Carlo S., and Mario V. (2004). A (sub)graph isomorphism algorithm for matching large graphs. *IEEE Trans. Pattern Anal. Mach. Intell.* , 26(10):1367–1372.
19. Medha A. (2015). *Left bit right: For SPARQL join queries with OPTIONAL patterns (left-outer-joins).* In: *Proc. of SIGMOD 2015*, pp.1793–1808.
20. Neumann T. and Weikum G. (2010). The RDF3X engine for scalable management of RDF data. *VLDB Journal*, 19(1):91–113.
21. Peng P., Zou L., Özsu M.T., Chen L., and Zhao D. (2016). Processing SPARQL queries over distributed RDF graphs. *VLDB J.*, 25(2): 243–268.
22. Pérez J., Arenas M., and Gutierrez C. (2009). Semantics and complexity of SPARQL. *ACM Trans. Database Syst.*, 34(3):30–43.
23. Picalausa F. and Vansummeren S. (2011). What are real sparql queries like? In: *Proc. of SWIM 2011*, article 7.
24. Prud'hommeaux E. and Seaborne A. (2008). SPARQL query language for RDF. *W3C Recommendation.*
25. Reinhard Pichler and Sebastian Skritek. Containment and equivalence of well-designed sparql. In: *Proc. of SIGMOD 2014*, pp. 39–50.
26. Saleem M., Ali M.I., Hogan A., Mehmood Q., and Ngomo A.C.N.(2015). LSQ: The linked SPARQL queries dataset. In: *Proc. of ISWC 2015*, pp. 261–269.
27. Schmidt M., Meier M., and Lausen G. (2010). Foundations of SPARQL query optimization. In: *Proc. of ICDT 2010*, pp. 4–33.
28. Shang H., Zhang Y., Lin X., and Yu J.X.(2008). Taming verification hardness: an efficient algorithm for testing subgraph isomorphism. In: *Proc. of VLDB 2008*, 1(1):364–375.

29. Song Z., Feng Z., Zhang X., Wang X., and Rao G.(2016). Efficient approximation of well-designed SPARQL queries. In: *Proc. of SemiBDMA at WAIM 2016*, LNCS 9998.
30. Swick R.R. (1998). Resource description framework (RDF) model and syntax specification. *W3C Recommendation*.
31. Ullmann J.R. (1976). An algorithm for subgraph isomorphism. *J. ACM*, 23(1):31–42.
32. Wang H. (2012). Efficient subgraph matching on billion node graphs. In: *Proc. of VLDB 2012*, 5: article 9.
33. Weiss C., Karras P., and Bernstein A. (2008). Hexastore: sextuple indexing for semantic web data management. In: *Proc. of VLDB 2008*, 1:1008–1019.
34. Yuan P., Liu P., Wu B., Jin H., Zhang W., and Liu L. (2013). Triplebit: A fast and compact system for large scale RDF data. In: *Proc. of VLDB 2013*, 6(7):517–528.
35. Zhang X., Feng Z., Wang X., Rao G., and Wu W. (2016). Context-free path queries on RDF graphs. In: *Proc. of ISWC 2016*, pp. 632–648.
36. Zhang X. and Van den Bussche J. (2015). On the power of SPARQL in expressing navigational queries. *Computer J.*, 58 (11): 2841-2851.
37. Zhang X. and Van den Bussche J. (2014). On the primitivity of operators in SPARQL. *Inf. Process. Lett.*, 114(9):480-485.
38. Zhang X. and Van den Bussche J. (2016). On the satisfiability problem for SPARQL patterns. *J. Artif. Intell. Res.* , 56:403-428.
39. Zhao P. and Han J. (2010). On graph query optimization in large networks. In: *Proc. of VLDB 2010*, 3(1-2):340–351.
40. Zou L., Mo J., Chen L., Özsu M.T., and Zhao D. (2011). gstore: answering sparql queries via subgraph matching. In: *Proc. of VLDB 2011*, 4(8):482–493.
41. Zou L., Özsu M.T., Chen L., Shen X., Huang R., and Zhao D. (2014). gStore: A graph-based sparql query engine. *VLDB J.*, 23(4):565–590.